Gaps and Opportunities for AI/ML Techniques in the EDA Domain

November 23, 2020

Artificial Intelligence and Machine Learning Special Interest Group

Joydip Das, Samsung Austin R&D Center (SARC)
Akhilesh Kumar, Ansys
Aparna Dey, Cadence Design Systems
Sashi Obilisetty, Google LLC
Prateek Bhansali, Intel Corporation
Srinivas Bodapati, Intel Corporation
Hui Fu, Intel Corporation
Anoop Saha, Mentor, a Siemens Business
W. Rhett Davis, NC State University
Ryan Carey, Qualcomm
Rajeev Jain, Qualcomm
Karthik V. Aadithya, Sandia National Laboratory
Leigh Anne Clevenger, Si2

Published by
Silicon Integration Initiative, Inc. (Si2™)
12335 Hymeadow Dr, Suite 450
Austin, TX 78750

Gaps and Opportunities for AI/ML Techniques in the EDA Domain

***Abstract*—Semiconductor industry awareness of AI/ML in EDA has progressed such that results from a user survey can drive industry-wide standards and improvements. Silicon Integration Initiative conducted such a survey in April 2020 identifying gaps, opportunities, and current practices for incorporation of AI/ML into the EDA domain. This paper presents and analyzes the findings of the survey. Areas explored include the current state of ML adoption in EDA at the respondents' organizations, and areas for potential improvement where respondents felt a certain level of dissatisfaction with the current state of ML availability and adoption within their organizations as well as their fields. Respondents identified their areas of interest related to EDA. The goal of this analysis is to understand the obstacles to adoption and potential improvements to AI and ML techniques in EDA. This will benefit EDA tool vendors and end-user engineers as well as individuals from academia, industrial design houses, and national laboratories. Future work by research and development groups can support these efforts through the development of a common AI/ML in EDA ecosystem.**

Keywords—artificial intelligence, EDA, electronic design automation, machine learning, semiconductor

I. Introduction

Artificial intelligence (AI) and machine learning (ML) have become fixtures in electronic design automation (EDA). Modern tool versions, specifications and publications exhibit this trend toward AI/ML integration. These efforts demonstrate improved tool behavior, particularly regarding the turnaround times and power-performance-areas (PPAs) of semiconductor designs.

The industry is rapidly approaching the next stage of AI/ML usage, leveraging the experiences of these recent endeavors to facilitate enhanced and extended implementation of AI/ML techniques. Silicon Integration Initiative (Si2), a research and development joint venture serving the semiconductor industry, conducted a survey in April 2020 identifying gaps, opportunities, and current practices for incorporation of AI/ML into the EDA domain. This paper presents and analyzes the findings of the survey.

Respondents included EDA tool vendors and end-user engineers, as well as individuals from academia, industrial design houses, and national laboratories. These respondents represent both analog and digital design domains. Survey questions were designed to capture respondent opinions on industry roadblocks and potential solutions.

There are three companion documents to this paper [1]:

- Complete survey results, including charts and tables
- Cross-tabulation results, comparing responses from specific areas of interest for each question (these include comparison of digital vs analog areas)
- Original survey questions

This paper is organized as follows: Section II describes the background of and motivation for the survey; Section III discusses related works; Section IV presents the limitations of the survey; Section V describes the survey questions; Section VI analyzes survey results by question; finally, Section VII draws conclusions and identifies avenues for further action.

II. Background and Motivation

EDA tools, techniques, and algorithms play a crucial role in modern semiconductor chip design. Every stage in the design of analog, digital, and mixed-signal systems—modeling, analysis, simulation, verification, and debugging—relies on the robust computational capabilities of cutting-edge EDA tools.

A wide array of ideas, concepts, and models drawn from several disciplines underpin these tools, including discrete math, graph theory, probability and statistics, continuous math (such as the theory of differential equations), matrix manipulation, and control theory. Bolstered by data availability and immense computational power, ML techniques have recently begun enabling significant theoretical and practical advancements in each of these disciplines, giving rise to powerful new algorithms and analysis techniques. Sizable ML-based projects are underway in a number of EDA-focused groups within academic institutions, national laboratories, EDA vendor companies, and chip design houses. There is a tremendous amount of investment by chip design houses, in particular, to build teams responsible for enabling efficient design and verification flows using AI/ML techniques, now perceived as crucial to maintaining competitiveness.

The confluence of these projects with the growing availability of AI/ML-capable systems forms the opportunity to create industry standard data models, frameworks, and best practices through collaboration among industry stakeholders. These efforts flatten the learning curve for such key topics as tools and techniques, libraries and platforms, cloud infrastructure and EDA-specific ML applications, as educating EDA professionals and their chip designer and system-integrator customers about core ML concepts becomes more viable and standardized. This project will also focus industry discussions on the drawbacks and limitations of ML methods such as over/under fitting, unjustified extrapolation, interpretability and confidence intervals, etc.

With technologies such as AI/ML, it is difficult to draw a clear boundary that allows for the required level of productive

knowledge sharing while protecting the intellectual property and private data of stakeholders. Without explicit delineation, any collaborative efforts would be at risk. Si2 is uniquely positioned to facilitate this exchange.

Understanding the current employment of AI/ML techniques in EDA production processes is a crucial first step toward crafting a collaborative framework to build, enable, debug, and deploy AI/ML-enabled solutions. A common understanding enables Si2, and the industry more broadly, to identify key areas where common standards, cooperation, and educational efforts would be most impactful. Such areas will form the greatest impetus for collaboration.

To this end (as mentioned in the introduction), an industry survey was conducted. Each participant's background and area(s) of interest in chip design and EDA tools were identified and probed by multiple-choice questions, such as:

- Which of these areas is most relevant to your work in the semiconductor industry?
- Which of these specific interests in EDA tool development do you share?

The survey evaluated two main areas:

- The current state of ML adoption in EDA at each participant's organization, using questions such as:
    - In what areas are you trying AI and ML methods?
    - What design data do you use or plan to use in your ML training and inference?
- Areas for potential improvement, where the participant felt a certain level of dissatisfaction with the current state of ML availability and adoption within their organization and field. Questions in this section included:
    - How satisfied are you with your company's progress in AI and ML, in product implementation, development, research, and individual research?
    - What resources would enable faster adoption of AI and ML for your team?
    - Have you experienced any of these roadblocks to EDA interoperability between EDA tool suppliers and users for AI and ML?

This work analyzes the results of the survey to inform the direction of future projects streamlining the adoption of AI/ML techniques in EDA to the benefit of all industry participants and stakeholders.

III. Related Work

While much EDA-focused AI/ML research likely remains unpublished due to concerns about protecting intellectual property and competitive advantage, several promising examples have nonetheless emerged that highlight the vital role of AI/ML in advancing EDA tools, techniques, and algorithms.

A large and growing body of work has surfaced, arising primarily from premier EDA conferences (e.g., DAC, ICCAD) and journals (e.g., TCAD). These reports guided the formation of this survey and may help inform future short and long-term AI/ML initiatives. This body of work has already aided Si2 in prioritizing development of common industry standards to foster interdisciplinary and industry-wide collaboration, enabling faster development, prototyping, benchmarking, and comparison of ML-based EDA techniques.

The existing literature on applying AI and ML to EDA is vast; thus, only a selection of key research initiatives and publications are covered here. This allows the reader to appreciate the breadth of projects and initiatives in this area, while providing a starting point for further investigation.

The Center for Advanced Electronics Through Machine Learning (CAEML) is an NSF-funded collaborative effort between the University of Illinois system, NC State University, and the Georgia Institute of Technology. CAEML is currently working to secure "EDA via ML" on a strong theoretical and practical foundation. The aim is to develop fast and accurate ML-based models for electronic systems and components at various levels of abstraction for large-scale micro-electronic design, simulation, and verification. This research initiative has resulted in a number of published works [2].

One example of AI/ML in EDA from academia was presented at ICCAD 2019 by UC San Diego [3]. The authors used machine learning models derived from electrical features of a given SOC floorplan and PDN. The machine learning methods predicted the updated static IR drop for each power node through different state changes. Their results achieved speedup compared to industry solutions.

EDA vendors have published papers in EDA conferences about ongoing projects and internal ML-focused efforts. One such paper, published in ASPDAC 2018, describes the role of computational techniques from ML (such as classification and regression) in shaping software supplier Synopsys's research in areas like functional and formal verification, yield modeling, and analog circuit performance characterization [4].

ML-based data-driven development of compact models—differential equations that describe the voltages, currents, and charges in individual devices on a semiconductor chip—is a key area of focus for Sandia National Laboratories. National labs like Sandia have ongoing research projects in ML-focused EDA [5].

Additionally, chip design houses often sponsor internal ML-focused efforts aimed at solving specific design problems encountered at advanced technology nodes. While many of these efforts go unpublished, some enter the proceedings of EDA conferences and workshops, and provide valuable insights into ongoing "ML for EDA" initiatives, such as a paper on parasitic extraction and estimation via ML published at DAC2020 by researchers at Intel [6].

Industry publications can provide general background on AI/ML concepts and trends, along with insights from subject matter experts. One example is the article "AI Begins to Reshape Chip Design" from Semiconductor Engineering [7].

The Special Interest Group has not found any survey on the application of AI/ML techniques in EDA domain; there are, however, surveys targeting other domains in the industry. The McKinsey Global Survey finds that most companies have observed measurable benefits from adoption of AI techniques [8]. According to the survey, AI adoption has increased in almost every industry from 2018 to 2019, with the largest increase being in retail (35%). The survey also reports much improvement in corporate ability "to scale impact, manage risks, and retain the workforce" as a result of incorporating AI. Additionally, the McKinsey survey finds that early adopters of AI techniques may have had advantages over their laggard competitors.

The authors are not aware of other AI/ML-in-EDA industry surveys with the number of responses (over 190) and wide distribution attained by this survey. There are smaller surveys from EDA companies to their customers, and surveys where AI/ML is the focus of only a few questions. Semiconductor industry awareness of AI/ML in EDA has reached a level where results from a user survey can and should drive industry-wide standards and improvements.

IV. Limitations of the Survey

The survey intended to cover the entire EDA ecosystem, from research, EDA, chip, and IP development to system and board design. The sample size of this survey is typical for this domain; however, it is worth noting that the majority of responses originated from EDA tool developers (49%), followed by chip designers (22%). Only 4% of respondents identified themselves as verification-focused engineers or developers. The uneven composition of survey respondents limits the capacity for well-rounded representation of the EDA ecosystem. EDA tool development, chip design, IP development (9%) and research (9%) were chosen as the four areas of expertise out of the initially designed eight domains. Areas in foundry, system design, board design and verification are included in overall response data but were dropped from the detailed cross-tabular analysis due to low participation.

Further questionnaires targeting foundry, verification, board, and system engineers would remedy the absence of related data, with focused technical questions to increase response rates. Communication with additional universities and government research laboratories would grant a wider picture of research successes and needs, as well as focused surveys for companies located outside of the United States. This survey does not distinguish between U.S. and international responses.

Information regarding the area-specific experience level and geo-location of respondents would offer greater insight into their individual perspectives, potentially allowing for more informed weighting or further categorization of results.

V. Survey Methodology and Contents

Semiconductor design is generally known to be the art of approximation of the physical world. This is highly defined by the designer expertise, best practices and heuristics learned over several years of experience.

The "AI/ML in EDA Gaps and Opportunities Survey" was developed in March and April of 2020 by the Si2 AI/ML in EDA SIG comprised of 20 Si2 member companies. The questions were refined by a six-company subgroup consisting of Ansys, IBM, Intel, Keysight Technologies, Samsung, and Synopsys. The survey was available to the public via the online platform SurveyMonkey from April 15 to May 18, 2020, and was publicized through use of the social media site LinkedIn, a blog post on the Si2 website, direct email contact from SIG group members, direct emails to the Si2 contact list, and various industry mailing lists, resulting in 193 responses.

Survey questions were designed to be efficient for respondents, with a target completion time of five minutes or less in total. Results indicate that 60% of respondents completed the survey in less than five minutes, and 82% of respondents required fewer than 10 minutes. Question formats ranged from multiple choice single-answer and multiple-choice multi-answer to questions using a Likert scale to gauge respondent opinions. Likert scale questions were formatted as a matrix to consolidate questions, and a weighted average has been calculated to approximate the prevailing opinion for each group. Survey responses were weighted 1 through 5—from "Very Unsatisfied" to "Very Satisfied"—and multiplied by the percentage of respondents that selected each option. For example, the weighted average of answers from respondents in production implementation for Question 10 is calculated below:

$$\begin{aligned} Weighted\ Average \\ &= (1 \times 0.1544) + (2 \times 0.1946) \\ &+ (3 \times 0.3221) + (4 \times 0.2081) \\ &+ (5 \times 0.1208) = 2.9463 \approx 2.95 \end{aligned}$$

A weighted average of 2.95 approximates to 3, the weight given for "Neither Satisfied nor Unsatisfied." Thus, we conclude that industry members involved in production implementation are generally ambivalent about their companies' overall progress with AI and ML.

An "Other" comment box was included in each question to reduce survey bias in case none of the listed options applied to the respondent. Question 1, the respondent's main area of interest, was the only required question, as this context was necessary to sort all results for analysis. The full text of the survey questions can be found in a companion document.

TABLE I. Question Categories

Category	**Corresponding Questions**
Primary Area of Interest	1
Secondary Area of Interest	2-9
Current AI/ML Methodology Gaps	10, 12
AI/ML Data	11, 12, 14
Reference AI/ML Flow	11,12, 16
Data Model	12, 16
Interoperability	12, 16
Education	13, 15
Software and Hardware	17, 18
Respondent Business Information	19, 20

The motivation for the survey was to gather relevant industry feedback from within and outside the SIG for the following areas: AI/ML methodology gaps, primary and secondary areas of respondent interest, company characteristics for each respondent, and the perceived industry need for the following: readily available training data, a consistent reference flow, a common data model, program interoperability, AI/ML-related education, and AI/ML-capable software and hardware. Table I links the topics with the corresponding survey questions.

VI. Results and Analysis

Questions 1-9: Primary and Secondary Areas of Interest

Question 1 identifies the semiconductor area most applicable to the respondent's work, defined as his or her primary area of interest. This question is integral to the survey analysis strategy, enabling all subsequent responses to be categorized and compared by respondent fields. The counts of responses per interest area reveal imperfect industry coverage and highlight some limitations of the survey.

Once Question 1 was completed, survey logic directed each respondent to an area-specific, multiple-choice, multi-answer question to refine the respondent's interests into specialty subfields. The survey logic for Questions 1-9 is summarized in Table II. Available answers for Questions 1-9 were proposed and approved by subject matter experts within the SIG. Responses to Questions 10-21 can thus be filtered by the response to Question 1 and the corresponding value(s) of Question 2-9. This information allows for the cross-sectional analysis of industry priorities conducted throughout this section. We will highlight some major findings here.

EDA tool development and chip design were the most common primary areas of interest, at 49% and 22% of respondents, respectively. Each remaining area—board design, foundry, IP development, research, system design, and verification—received less than 10% of responses; however, Questions 2-9 reveal overlap between the primary areas of interest. For instance, while only 4% of respondents identified verification as their primary area of interest, more than 34% of EDA tool developers have identified verification among their secondary areas of interest. Furthermore, survey results indicate a healthy mix of respondents from both digital and analog/mixed-signal design domains.

TABLE II. Survey Logic for Questions 1-9

Primary Area of Interest (Question 1)	Corresponding Secondary Interest Question
Chip Design	Question 2
IP Development	Question 3
EDA Tool Development	Question 4
Foundry	Question 5
System Design	Question 6
Board Design	Question 7
Research	Question 8
Verification	Question 9
Other (please specify)	None

Question 10: How satisfied are you with your company's progress in AI and ML?

Question 10 gauges adoption of AI/ML and the communication of organizational strategies around AI; Table III captures the results. The weighted averages for Question 10, described in Section III, indicate that the semiconductor industry may be on the path toward meaningful adoption, but further improvements and education are required.

This result suggests that the semiconductor industry remains in the early phases of real production deployment, and researchers and engineers are still exploring for AI/ML development. The dominance of the neutral opinion supports this conclusion. The fact that small scale deployment is more feasible for individual and small group efforts than for industry titans may account for the satisfied responses. Potential explanations for respondent dissatisfaction include the inherent challenges of large enterprises, where meaningful change requires investment, ROI analysis, and management buy-in, as well as detailed plans to persuade the end user community. Lack of standardization, separate setups for ML environment and design environments, and compelling education for end users may also contribute to the dissatisfaction observed in responses.

In a more detailed analysis, Table IV cross-tabulates responses by secondary areas of interest. A significant portion of respondents across various subfields indicate partial or complete satisfaction with organizational commitments to AI/ML. Furthermore, anecdotal evidence from respondent comments suggest that a large segment of end users have begun learning about AI/ML of their own initiative, which bodes well for AI/ML implementation at the organizational level.

Question 11: In what areas are you trying AI/ML methods?

Respondents were asked about the specific areas in which AI/ML techniques are being tested within their organization, if any. The responses are presented in Table V.

Simulation (41%), place and route (38%), and compute efficiency (33%) stand out as early areas in which ML is being applied. These fields are ripe for ML implementation, as they pose significant bottlenecks for the iterative design process and have readily available data for training. Crucially, these scenarios can be framed as so-called "supervised learning problems," for which state-of-the-art ML is quite advanced, resulting in meaningful positive applications such as those explained below.

TABLE III. How satisfied are you with your company's progress in AI and ML?

	Very Unsatisfied	Unsatisfied	Neither Satisfied Nor Unsatisfied	Somewhat Satisfied	Very Satisfied	Total	Weighted Average
Product Development	15.44%	19.46%	32.21%	20.81%	12.08%	149	2.95
Implementation	12.00%	21.33%	30.67%	22.00%	14.00%	150	3.05
Research	9.46%	18.92%	27.03%	25.00%	19.59%	148	3.26
Individual Research	10.88%	16.33%	34.69%	20.41%	17.69%	147	3.18
						Answered	**153**
						Skipped	**40**

TABLE IV. CROSS-TABULATION OF RESPONSES TO QUESTION 10.

	Very Unsatisfied	Somewhat Unsatisfied	Neither Satisfied Nor Unsatisfied	Somewhat Satisfied	Very Satisfied	Weighted Average
	Production Implementation					
Tool Developers	16%	16%	32%	25%	11%	2.99
Chip Designers	14%	11%	33%	22%	19%	3.22
Researchers	11%	56%	22%	0%	11%	2.44
IP Developers	14%	36%	29%	14%	7%	2.64
All Respondents	15%	19%	32%	21%	12%	2.95
	Development					
	Very Unsatisfied	Somewhat Unsatisfied	Neither Satisfied Nor Unsatisfied	Somewhat Satisfied	Very Satisfied	Weighted Average
Tool Developers	15%	18%	27%	27%	14%	3.07
Chip Designers	6%	25%	33%	17%	19%	3.19
Researchers	11%	11%	44%	22%	11%	3.11
IP Developers	13%	27%	33%	20%	7%	2.80
All Respondents	12%	21%	31%	22%	14%	3.05
	Research					
	Very Unsatisfied	Somewhat Unsatisfied	Neither Satisfied Nor Unsatisfied	Somewhat Satisfied	Very Satisfied	Weighted Average
Tool Developers	12%	18%	26%	22%	22%	3.23
Chip Designers	0%	31%	33%	19%	17%	3.22
Researchers	10%	0%	30%	30%	30%	3.70
IP Developers	14%	14%	14%	50%	7%	3.21
All Respondents	9%	19%	27%	25%	20%	3.26
	Individual Research					
	Very Unsatisfied	Somewhat Unsatisfied	Neither Satisfied Nor Unsatisfied	Somewhat Satisfied	Very Satisfied	Weighted Average
Tool Developers	13%	15%	32%	25%	15%	3.15
Chip Designers	8%	22%	36%	11%	22%	3.17
Researchers	10%	0%	40%	20%	30%	3.60
IP Developers	14%	7%	43%	21%	14%	3.14
All Respondents	11%	16%	35%	20%	18%	3.18

- Place and route iterations arise from weak congestion and timing prediction and poor correlation between synthesis, placement and routing results. ML-based data-driven models are key to push PPA and reduce overdesign in the post-Moore era.
- Simulations to sweep design parameters can be very time consuming, and the eventual simulation output for one parameter set is often related to output for another. ML models can capture this relationship, reducing the number of simulations run and lessening the bottleneck effect.
- Compute efficiency requires optimizing compute usage for various EDA applications. In most cases, data from past projects is abundant and available for usage and optimization training.

These three areas share access to relevant data required to model supervised learning, which, when combined with their importance to the design process, may explain the high satisfaction ratings received.

Table VI, which categorizes the responses in accordance with each respondent's primary area of interest, reveals that tool developers and researchers are most likely to apply these techniques to simulation, while chip designers are more likely to employ AI/ML methods during place and route. We also find

TABLE V. IN WHAT AREAS ARE YOU TRYING AI/ML METHODS?

Answer Choices	Responses
Simulation	41.41%
Place and Route	37.50%
Compute Efficiency	32.81%
Verification and Debug	30.47%
Design Rule Checking	22.66%
Regression Analysis	22.66%
Root Cause Analysis	19.53%
Other (please specify)	17.97%
Building Standard Cells	14.84%
Improving Coverage	14.06%
Emulation	7.03%
Formal Property Verification	6.25%
	Answered 128
	Skipped 65

that researchers are mainly focusing on the improvement of compute efficiency. This is to be expected, since many AI/ML applications aim to improve the computational efficiency of classical methods. Formal verification and emulation stand out as areas with minimal current focus, though this is likely an artifact of minimal survey participation from that domain.

While many tool developers are applying AI/ML to simulation, the majority work in analog and mixed-signal design. It is worth noting that more than one third of survey respondents identify as working in analog and mixed-signal design, where simulation tools play a major role. Likewise, the implied dependence of chip designers on AI/ML techniques for place and route may arise from the composition of survey respondents, 70% of whom expressed interest in digital design and physical design and analysis.

In areas such as verification and debug, the lack of sufficient data to train a model may contribute to the low rating apparent in the survey results. For ML to be successful in these areas, new techniques that can learn from limited or synthetically generated data are required. Management and business decision-makers seem to believe that their companies are most likely to benefit from AI/ML methods applied to verification and debug. Perhaps this belief is due to verification and debug consuming a large portion of the design cycle, and accelerating this process would, in turn, accelerate time to market. Note that there is a possibility for overlap between verification and debug tools and those of other subfields; for example, a simulation-based verification tool can be classified as either a simulation tool or a verification and debug tool.

TABLE VI. CROSS-TABULATION OF RESPONSES TO QUESTION 11.

	Building Standard Cells	Simulation	Place & Route	Design Rule Checking	Improving Coverage	Verification & Debug
Tool Developers	15%	53%	27%	25%	15%	32%
Chip Designers	15%	21%	64%	24%	18%	33%
Researchers	11%	56%	33%	11%	0%	11%
IP Developers	23%	46%	31%	23%	15%	38%
All Respondents	15%	41%	38%	23%	14%	30%

continued...	Root Cause Analysis	Formal Property Verification	Regression Analysis	Compute Efficiency	Emulation	Other
Tool Developers	22%	10%	24%	34%	8%	22%
Chip Designers	21%	3%	21%	18%	3%	9%
Researchers	11%	0%	44%	67%	0%	11%
IP Developers	23%	8%	23%	38%	15%	8%
All Respondents	20%	6%	23%	33%	7%	18%

Question 12: What is important to enable faster implementation of AI and ML EDA solutions?

Question 12 touches important vectors for enabling faster implementation of AI/ML EDA solutions like methodology gaps, data availability, example flows, and interoperability, as well as the need for a common data model. The focus here is on analysis of high-interest areas. Tables VII and VIII present the responses to Question 12, categorized by the primary interest areas and primary organizational roles of the respondents, respectively.

The majority of respondents indicated that readily available training data would accelerate implementation of AI/ML techniques, with a weighted average for the question of 4.46. This result implies that many AI/ML practitioners do not have access to relevant data for training. Data that would be helpful is generally binary in nature, part of unstructured log files, or internal to EDA tools. A related conclusion is the lack of a common data model for storing and passing data around an EDA tool stack, with 69% of respondents describing a common data model as either "Important" or "Very Important."

Survey respondents gave a mean importance score of 4.17 to "appropriate" algorithms for data in EDA. Useful data such as time-series data and graphs may be derived in nature for EDA [9]. In this context, specialized algorithms that scale and perform well for AI/ML in EDA may be needed, such as GCN/GNN.

TABLE VII. WHAT IS IMPORTANT TO ENABLE FASTER IMPLEMENTATION OF AI AND ML EDA SOLUTIONS?

	NOT IMPORTANT	SOMEWHAT IMPORTANT	NEITHER IMPORTANT NOR UNIMPORTANT	IMPORTANT	VERY IMPORTANT	TOTAL	WEIGHTED AVERAGE
Common Data Model	3.00% 3	8.00% 8	11.00% 11	49.00% 49	29.00% 29	100	3.93
Higher Level Model Abstractions	4.08% 4	6.12% 6	19.39% 19	51.02% 50	19.39% 19	98	3.76
Training Data Availability	0.00% 0	2.86% 3	8.57% 9	28.57% 30	60.00% 63	105	4.46
Appropriate Algorithms for Data	0.97% 1	2.91% 3	9.71% 10	51.46% 53	34.95% 36	103	4.17
Reference Flow Availability	0.97% 1	1.94% 2	17.48% 18	52.43% 54	27.18% 28	103	4.03
Analysis Optimization Engine	3.03% 3	4.04% 4	19.19% 19	48.48% 48	25.25% 25	99	3.89
Models for Software Settings	4.90% 5	6.86% 7	32.35% 33	45.10% 46	10.78% 11	102	3.50
Methodology to Reuse Data	4.04% 4	3.03% 3	16.16% 16	46.46% 46	30.30% 30	99	3.96
Compute Capacity	1.94% 2	6.80% 7	16.50% 17	45.63% 47	29.13% 30	103	3.93
Licensing Methodology for ML Model Reuse	8.00% 8	14.00% 14	23.00% 23	39.00% 39	16.00% 16	100	3.41

More than 75% of respondents rated a uniform reference flow as "Important" or "Very Important." This standard format would reduce ramp-up time for ML practitioners in EDA. ML model generation experience does not translate well across domains like NLP, computer vision, etc.; models from one of these disciplines must be specifically tuned for EDA, due to the inherent differences in numerical or categorial data between EDA and other established fields.

Methodology to reuse data and compute capacity also scored highly, with three-quarters of respondents describing each as "Important" or "Very Important." Given the dependence of ML on past data, and the necessity of compute clusters for model training, this is an intuitive result.

Question 13: Inherent in any ML analysis is a percentage of uncertainty. What would make AI and ML solutions trustworthy and successful?

Given the imperfect nature of AI/ML, respondents were asked to gauge the impact of various mitigation methods on their use of ML-driven algorithms. Table IX presents their responses and Table X cross-tabulates the results based on the primary areas of interest of the respondents.

Machine learning model predictions contain some degree of inherent inaccuracy, since the model is extrapolating what it learns from known data to unknown data; by contrast, current commercial EDA tools conduct an exact computation of the desired metric for each dataset. This computation would be the benchmark for comparison against the ML model prediction. Given these inevitable discrepancies between the prediction and the EDA tool computation, designers expressed understandable wariness of using ML tools without conventional methods as a check.

Respondents strongly gravitated toward a use model in which AI/ML is used in concert with standard methods, applying AI/ML methods to augment the early stages of design and traditional algorithms for final approval and signoff of designs. Though a majority of respondents were in favor of clear definitions for applications where different but correct solutions are acceptable (57%), and clear definitions of tradeoffs with exhaustive analysis (61%), the most trusted use of AI/ML was limited to the early stages of design (64%). Even this approach only garnered 64% support, perhaps due to designer fears that ML inaccuracies may lead to additional iterations at signoff if the traditional tools discover too many design violations. One method to mitigate this risk and enhance the palatability of ML tools would be to include an uncertainty or confidence prediction with every ML prediction, and re-train the ML model or resort to traditional tools if uncertainty falls below a

TABLE VIII. CROSS-TABULATION OF RESPONSES TO QUESTION 12.

Weighted Averages	Common Data Model	Higher Level Model Abstractions	Training Data Availability	Appropriate Algorithms for Data	Reference Flow Availability	Analysis Optimization Engine	Models for Software Settings	Methodology to Reuse Data	Compute Capacity	Licensing Methodology for ML Model Reuse
Tool Developers	3.88	3.60	4.55	4.19	4.04	3.88	3.43	3.90	3.82	3.42
Chip Designers	3.95	3.87	4.55	4.09	3.87	3.87	3.41	3.77	4.09	3.09
Researchers	3.83	3.83	4.50	4.17	4.20	4.17	3.83	4.00	4.00	4.17
IP Developers	4.50	4.20	4.09	4.33	4.30	3.67	3.80	4.60	4.30	3.80
All Respondents	3.93	3.76	4.46	4.17	4.03	3.89	3.50	3.96	3.93	3.41

TABLE IX. WHAT WOULD MAKE AI AND ML SOLUTIONS TRUSTWORTHY AND SUCCESSFUL?

Answer Choices	Responses
Use AI/ML in Early Stages of Design, and Traditional Algorithms for Final Sign-Off	64.42%
Clearly Define Acceptable Trade-offs to Exhaustive (100%) Analysis, for example 99% or 90%	60.58%
Clearly Define Applications Where Different, Correct Solutions Are Acceptable	56.73%
Use Deterministic AI without ML	17.31%
Other (please specify)	5.77%
	Skipped 89
	Answered 104

TABLE X. CROSS-TABULATION OF RESPONSES TO QUESTION 13.

	Clearly Define Applications Where Different, Correct Solutions Are Acceptable	Clearly Define Acceptable Trade-offs to Exhaustive (100%) Analysis, for example 99% or 90%	Use AI/ML in Early Stages of Design, and Traditional Algorithms for Final Sign-Off	Use Deterministic AI without ML	Other
Tool Developers	58%	62%	62%	18%	4%
Chip Designers	57%	52%	74%	22%	4%
Researchers	80%	80%	60%	20%	20%
IP Developers	55%	55%	73%	0%	0%
All Respondents	57%	61%	64%	17%	6%

predetermined confidence level.

Defining acceptable trade-offs, the second-most trusted option, is not mutually exclusive with the above actions. Design KPIs that can be adversely affected by the inaccuracies of the ML prediction must be defined and correlated with model uncertainty. This method would compute an accuracy target for which KPIs would not be disrupted.

This question shows how respondents view the importance of incorporating knowledge expertise in addition to EDA data. As shown in Table X, between 50% and 64% of respondents indicated that they need guidance on where in the design flow to best leverage AI/ML solutions, with 80% of researchers recognizing the importance in understanding acceptable trade-offs. Different applications (i.e. optimization steps) will require different levels of accuracy, and it is therefore not realistic to give a single acceptability criterion for all steps. At each stage of a typical EDA design flow, more data is available than at the last, which allows for broader AI/ML usage. Responses may be skewed by competing measures of acceptability, even among respondents from the same application domain. Future work may provide domain-specific and/or generic metrics for acceptability to clarify these choices for users.

Question 14: What design data do you use or plan to use in your ML training and inference? (choose all that apply)

AI/ML techniques typically require a large set of data for successful training. Question 14 asked respondents about the current or planned datasets in use for ML training. Table XI presents the results. Simulation data (60%) appears to be used significantly more often than the other forms of data listed. The second tier of responses—layout data, place and route data, timing data and power data—represent forty to fifty percent of respondents who are using one or more of these data sets for ML training and inference.

TABLE XI. WHAT DESIGN DATA DO YOU USE OR PLAN TO USE IN YOUR ML TRAINING AND INFERENCE?

Answer Choices	Responses
Simulation	60.20%
Layout	47.96%
Place and Route	45.92%
Timing	45.92%
Power	40.82%
Verification	33.67%
Design Rules	31.63%
Standard Cells	29.59%
Noise	22.45%
Manufacturing	15.31%
Other (please specify)	9.18%
	Skipped 95
	Answered 98

Table XII presents cross tabular data that breaks down demographic information for these responses. As observed above, simulation data is widely used for ML applications; however, this table presents some interesting details. For instance, more than 70% of chip designers use or anticipate using place and route or timing data for their ML applications, while only roughly 40% of the other respondent groups use such data for their ML applications. Researchers, in particular, heavily depend on simulation data for their ML applications.

Table XIII presents data comparing results between respondents from the digital and analog domains. Simulation results are more commonly used for analog domain (80%+), while timing information is more commonly used for digital domain (72%+). Since timing is a prominent quality metric for digital design, higher usage among digital designers than analog designers is to be expected.

The results presented in Tables XII and XIII are logical and intuitive. Simulation data are more readily available than other data forms, and the results naturally indicate a high level of usage. Since researchers—especially those in academia—have reduced access to manufacturing data, noise data, and realistic place and route data, they rely heavily on simulation data. Perhaps secure knowledge diffusion methodologies could help domain experts make datasets more readily available and viable in other domains.

Additional data types used for AI/ML from the write-in responses to Question 14 include SRAM, PPA data, CDC checks, configuration, I-V and C-V data from transistor curve tracer measurement, analog design centering, and user training.

Further efforts could investigate the origins of data that respondents are using or plan to use, and whether lack of availability of a particular data type is responsible for its low implementation. While this question focused on design data, a similar question could ascertain industry opinions about derived data from software input and output, or metadata, such as that arising from customer support or Q&A systems.

TABLE XII. CROSS-TABULATION OF RESPONSES TO QUESTION 14

	Standard Cells	Simulation	Layout	Place & Route	Timing	Noise	Power	Design Rules	Verification	Manufacturing	Other
Tool Developers	31%	60%	52%	40%	44%	20.83	46%	29%	33%	10%	8%
Chip Designers	33%	57%	52%	71%	71%	28.57	57%	24%	33%	29%	10%
Researchers	20%	80%	40%	40%	20%	0.00	20%	40%	20%	0%	0%
IP Developers	27%	64%	18%	36%	45%	27.27	36%	36%	36%	9%	18%
All Respondents	30%	60%	48%	46%	46%	22.45	41%	32%	34%	15%	9%

TABLE XIII. RESPONSES TO QUESTION 14 CATEGORIZED BY ANALOG AND DIGITAL DOMAINS

Analog vs. Digital	Standard Cells	Simulation	Layout	Place & Route	Timing	Noise	Power	Design Rules	Verification	Manufacturing	Other
Analog/Mixed-Signal Design											
Tool Developers	40%	80%	63%	40%	43%	23%	47%	27%	27%	3%	10%
Chip Designers	57%	86%	71%	71%	57%	29%	57%	43%	43%	43%	0%
Digital Design											
Tool Developers	44%	50%	50%	61%	72%	33.33	66.67	33.33	27.78	5.56	6%
Chip Designers	36%	64%	57%	57%	79%	29%	64%	36%	43%	36%	14%

Question 15: What resources would enable faster adoption of AI and ML for your team?

The most notable observation from the responses to Question 15, presented in Table XIV, is that an overwhelming majority of respondents believe a uniform EDA reference flow and online AI/ML classes on EDA applications would speed up the adoption of AI/ML in their teams. An example reference flow would clarify the tools involved and how AI/ML can be used in such flows.

Slightly less importance is given to a video overview advocating AI/ML in EDA, perhaps because more respondents are already becoming aware of beneficial AI/ML applications in EDA. This implies that the aforementioned reference flow, with an example showing the application of AI/ML in EDA, is especially crucial for industry adoption.

Additionally, while many online classes and courses exist for AI/ML and their applications, there appears to be a dearth of online classes teaching AI/ML as applied explicitly to EDA, despite the fact that there are a growing number of AI/ML applications in EDA and publication of papers in this area. A joint initiative between industry and academia to create high quality online classes could bridge this gap.

More than 80% of respondents favor EDA-specific online classes and reference implementation for faster adoption of AI/ML in their own teams. The responses to this question indicate a clear gap between the resources publicly available online from existing AI/ML courses and those targeted toward semiconductor engineers. With online courses becoming more popular, a variety of learning resources are readily available on platforms like Coursera; however, most of the AI/ML applications presented in these courses do not resonate well with analog designers, digital designers, or EDA tool developers [10]. For online training to be helpful to EDA stakeholders, it must include industry-specific applications like placement, routing, and circuit simulation, all of which are currently unavailable.

TABLE XIV. WHAT RESOURCES WOULD ENABLE FASTER ADOPTION OF AI AND ML FOR YOUR TEAM?

Answer Choices	Responses
On-Line Classes on Building and/or Using AI and ML in EDA Applications	79.61%
Downloadable or Cloud-Based AI/ML in EDA Reference Flow with Example Data and Algorithms	77.67%
Video on Understanding AI/ML and What It Can Do for EDA	55.34%
Other (please specify)	5.83%
	Skipped 90
	Answered 103

Question 16: Have you experienced any of these roadblocks to EDA interoperability between EDA tool suppliers and users for AI and ML?

As shown in Table XV, chip designers are noticeably concerned about a lack of training data. While design teams have generations of design and derived data available, they lack a common way to organize this data for AI/ML training and inference, making it difficult or impossible to take advantage of the available data for ML. This is a substantial roadblock to collaboration with tool developers both internal to and external of their respective companies.

As shown in Table XVI, respondents from all fields are concerned about the lack of a common data model. Although each development group has its own proprietary data model, combining the common training, inference data organization, and design flow for demonstration is important for all respondents, including those from digital, analog, and verification domains.

This question has a terminological limitation, as interoperability issues may exist not only between the tools from different vendors, but also between the tools from the same vendor. The question does not allow us to ascertain which aspect of the issue the respondents considered while answering this question, and may, in fact, undervalue the challenge interoperability issues pose to AI/ML implementation.

TABLE XV. HAVE YOU EXPERIENCED ANY OF THESE ROADBLOCKS TO EDA INTEROPERABILITY BETWEEN EDA TOOL SUPPLIERS AND USERS FOR AI AND ML?

Answer Choices	Responses
Lack of Training Data	65.59%
No Common Reference Flow	56.99%
No Common Data Model	49.46%
No licensing methodology for existing ML models	22.58%
Other (please specify)	9.68%
I Have No Interoperability Problems	8.60%
	Skipped 100
	Answered 93

TABLE XVI. RESPONSES TO QUESTION 16 BY RESPONDENT INTEREST AREA.

	No Common Data Model	Lack of Training Data	No Common Reference Flow	No Licensing Methodology for Existing ML Methods	I have no Interoperability Problems	Other
Tool Developers	51%	71%	56%	22%	11%	7%
Chip Designers	52%	71%	52%	24%	5%	5%
Researchers	60%	60%	40%	20%	20%	0%
IP Developers	60%	70%	90%	30%	0%	0%
All Respondents	49%	66%	57%	23%	9%	10%

The existence of a common data model supported by a common framework may help address these interoperability concerns; additionally, end users must work with vendors to ensure access to the derived data from flow steps. As a side benefit, these data will enable end users to provide enhanced feedback to vendors and developers.

Question 17: How satisfied or dissatisfied are you with your AI and ML hardware platforms?

Table XVII displays the relative satisfaction levels of respondents with their organizations' current AI/ML hardware implementations.

It is surprising to see such a strong positive response for TPU in this survey, given that there are few EDA tools and minimal IC design workloads running on TPUs. This result may reveal more about ML training workloads generally than EDA-specific ML inference workloads. Cloud-based solutions present a similar case, given the small portion of EDA workloads presently in the cloud. Both options have a notable percentage of respondents selecting "Not Applicable" (32% and 20%, respectively). These abstentions could explain the relatively high satisfaction level with GPU/TPU, though only a small population—engineers directly conducting ML training—stand to benefit.

These results also suggest an opportunity for EDA companies to utilize the power of GPU and TPU for workloads both related and unrelated to ML; however, that may require a major software architecture overhaul for application to current EDA solutions.

TABLE XVII. HOW SATISFIED OR DISSATISFIED ARE YOU WITH YOUR AI AND ML HARDWARE PLATFORMS?

	Very Dissatisfied	Dissatisfied	Neither Satisfied Nor Dissatisfied	Satisfied	Very Satisfied	Not Applicable	Total	Weighted Average
Cpu	2.20%	8.79%	35.16%	29.67%	20.88%	3.30%	91	3.68
Gpu	0.00%	11.11%	37.78%	28.89%	12.22%	10.00%	90	3.72
Tpu	0.00%	6.90%	37.93%	16.09%	6.90%	32.18%	87	4.2
Cloud	0.00%	10.34%	34.48%	25.29%	9.20%	20.69%	87	3.95
Other (please specify)							1	
							Answered	**93**
							Skipped	**100**

Question 18: How satisfied or dissatisfied are you with your AI and ML software platforms?

There has been tremendous innovation in AI frameworks in the recent past; TensorFlow, PyTorch, MxNet come to mind. To better understand the suitability of said frameworks for EDA, respondents were asked about their satisfaction with current AI/ML software platforms. Respondents were generally satisfied with the state of software platforms for enabling AI/ML, as seen in Table XVIII.

Over 75% of respondents reported favorably on the current state of their organizations' software platforms. Both training and inference environments were considered, and were deemed largely satisfactory across all industry segments surveyed. It appears that AI and ML software platforms are not roadblocks to AI/ML implementation, so no new work in this area is required to enable EDA solutions. Individual companies should continue to monitor new computer science innovations that are appropriate for their data and projects.

Question 19: Your company's annual revenues.

Table XIX presents the approximate annual revenues of survey respondent companies and organizations. The majority (58%) of survey participants who responded to this question belonged to companies with annual revenues of $1 billion or more. While this information provides interesting context, it is not fully representative of the survey participants, since 52% of participants declined to answer this question and another 8% of participants were reportedly unsure of their company's revenue. A future survey could target start-up and mid-sized companies to discern whether their needs are consistent with those of the majority of the respondents in this survey.

TABLE XVIII. HOW SATISFIED OR DISSATISFIED ARE YOU WITH YOUR AI AND ML SOFTWARE PLATFORMS?

	Model Development & Training Environment							Inference Environment / Algorithms						
	Very Dissatisfied	Dissatisfied	Neither Satisfied Nor Dissatisfied	Satisfied	Very Satisfied	Not Applicable	Weighted Average	Very Dissatisfied	Dissatisfied	Neither Satisfied Nor Dissatisfied	Satisfied	Very Satisfied	Not Applicable	Weighted Average
Tool Developers	0%	16%	26%	39%	16%	3%	3.63	0%	13%	29%	39%	13%	5%	3.68
Chip Designers	10%	14%	24%	38%	10%	5%	3.38	10%	5%	43%	29%	10%	5%	3.38
Researchers	0%	20%	40%	20%	20%	0%	3.40	0%	0%	60%	20%	0%	20%	3.80
IP Developers	0%	20%	50%	30%	0%	0%	3.10	0%	10%	50%	40%	0%	0%	3.30
All Respondents	2%	14%	33%	36%	10%	5%	3.52	2%	8%	38%	34%	10%	7%	3.63

	Applying Models for Multiple Uses							Big Data / Data Engineering							Other
	Very Dissatisfied	Dissatisfied	Neither Satisfied Nor Dissatisfied	Satisfied	Very Satisfied	Not Applicable	Weighted Average	Very Dissatisfied	Dissatisfied	Neither Satisfied Nor Dissatisfied	Satisfied	Very Satisfied	Not Applicable	Weighted Average	Other
Tool Developers	0%	16%	30%	35%	14%	5%	3.62	0%	16%	38%	24%	11%	11%	3.62	0%
Chip Designers	14%	10%	48%	14%	5%	10%	3.14	10%	15%	40%	25%	5%	5%	3.15	0%
Researchers	20%	20%	20%	20%	0%	20%	3.20	0%	0%	60%	20%	0%	20%	3.80	0%
IP Developers	0%	30%	50%	20%	0%	0%	2.90	0%	20%	70%	10%	0%	0%	2.90	0%
All Respondents	5%	16%	37%	24%	9%	8%	3.42	2%	14%	46%	21%	7%	9%	3.45	0%

TABLE XIX. YOUR COMPANY'S ANNUAL REVENUES.

Answer Choices	Responses
$1 billion - $10 billion	38.71%
More than $10 billion	18.28%
Unsure	16.13%
Less than $1 million	8.60%
$1 million - $5 million	6.45%
$100 million - $500 million	4.30%
$500 million - $1 billion	3.23%
$10 million - $50 million	2.15%
$5 million - $10 million	1.08%
$50 million - $100 million	1.08%
	Skipped 100
	Answered 93

Question 20: How did you hear about the survey?

A multi-vector contact approach accrued over 190 survey responses, a sizable pool for an EDA-focused survey. Table XX presents the breakdown of responses. The large number of "Si2 Email" responses is to be expected, since much effort promoting the survey went into direct email contact. Responses for "Colleague" and "LinkedIn" indicate enthusiasm for the project, and highlight the biggest opportunities for improvement. These results suggest that by encouraging team members who are active on LinkedIn to share announcements, and possibly by using LinkedIn ads, companies may significantly increase the exposure of future projects and surveys.

VII. FUTURE WORK AND CONCLUSIONS

Semiconductor industry engineers, designers, tool developers, and managers are pursuing AI/ML in a wide variety of areas, as shown in Questions 11 and 14. A common AI/ML in EDA infrastructure to support software tools and metrics would remove roadblocks identified by survey respondents and enable faster adoption of AI/ML in EDA, to the benefit of the industry. Development of a common AI/ML in EDA ecosystem can be achieved by research and development joint ventures (such as the survey sponsor, Si2), individual universities, or other development groups committed to open standards and solutions.

TABLE XX. HOW DID YOU HEAR ABOUT THE SURVEY?

Answer Choices	Responses
Si2 Email	57.95%
Colleague	31.82%
LinkedIn	20.45%
Other (please specify)	2.27%
Trade Publication	1.14%
	Skipped 105
	Answered 88

To address the glaring lack of classified and labeled training data identified in Questions 12 and 16, a standard common data model could be designed to represent derived data with classification and labeling for AI/ML pertaining to IC design, with a defined interface to design data (such as that of OpenAccess, for example). Additionally, industry members must conduct efforts to facilitate cross-company, cross-platform learning, without the need for design and EDA companies to disclose their IP. Building an ImageNet-type computer vision open database is not realistic for IC design and EDA tool vendors; however, federated learning via distrusted training could be explored to tap the data pools at individual companies.

Question 13 highlights roadblocks to confidence in AI/ML flows. A common verification and debug process could be defined for use with ML flows. Quality and confidence metrics can be developed, and software designed to output metrics to provide guidance to the user without exposing vendor IP. A checklist can be developed to assist projects in evaluating ML risk. User education in applying these metrics would be valuable for all respondents. The goal of establishing a common process is to provide transparency on decision-making within the ML model, and to give the end user the opportunity to judge the risk level. This would not be a performance benchmark like MLperf for AI/ML hardware, but rather, a set of common industry quality metrics to guide ML in EDA implementation. For example, to generate inference confidence intervals to evaluate precision of a result, the end user requires a means of finding and understanding the location of their specific use case data within the overall training dataset from the vendor. A data assessment indictor for end user data, providing a relative position for the overall training dataset, is also important.

Comments to Question 11 remarked on the application of AI/ML techniques to customer support scenarios, as well as customer-facing documentation. This is a rich source of data which could by analyzed by established natural language processing techniques.

Questions 12 and 16 discuss reusing ML models. Such models can include functions that provide feedback to the user on the logic behind a certain prediction, without exposing the internal model code. Metadata can be developed to confirm that consistent characterization is used for training and inference data sets.

Question 17 inquired about respondents' current cloud environments. Future AI/ML work could focus on ML metadata and data collection opportunities which are unique to the cloud, and EDA can learn from industries like automotive and medical in this area.

Survey respondents confirmed that the lack of classified and labeled industry-level training and inference data forms a significant roadblock to advancing AI/ML in EDA from research to production. Responses revealed the range of potential applications and types of data that are being tried with AI/ML. Respondents have a need for metrics and tools to evaluate and share models securely, and for education on the best practices for applying AI/ML in EDA.

The survey authors look forward to moving ahead on a common infrastructure to enable faster adoption of AI/ML in EDA.

REFERENCES

[1] J. Das, et al., "Gaps and Opportunities for AI/ML Techniques in the EDA Domain", Nov. 2020, Si2 Website.

[2] Grainger College of Engineering. "Center Publications." Center for Advanced Electronics through Machine Learning (CAEML). https://publish.illinois.edu/advancedelectronics/center-publications/ (accessed Oct. 19, 2020).

[3] C-T Ho, A.B. Kahng, "IncPIRD: Fast Learning-Based Prediction of Incremental IR Drop", ICCAD-2019. https://vlsicad.ucsd.edu/Publications/Conferences/375/c375.pdf (accessed Oct. 15, 2020).

[4] M. Pandey, "Machine learning and systems for building the next generation of EDA tools," in ASPDAC '18: Proceedings of the 23rd Asia and South Pacific Design Automation Conference, Jeju, Republic of Korea, Jan. 2018, pp. 411-415.

[5] K. Aadithya et al., "Development, demonstration, and validation of data-driven compact diode models for circuit simulation and analysis," Jan. 2020. https://arxiv.org/abs/2001.01699

[6] P. Bhansali et al, "MLParest: machine learning based parasitic estimation for custom circuit design," Design Automation Conference, Virtual, June 2020.

[7] E. Sperling, "AI begins to reshape chip design," Semiconductor Engineering, Nov. 1, 2018. https://semiengineering.com/ai-begins-to-reshape-chip-design/ (accessed Oct. 19, 2020).

[8] McKinsey & Company, "Global AI survey: AI proves its worth, but few scale impact," Nov. 22, 2019. https://www.mckinsey.com/featured-insights/artificial-intelligence/global-ai-survey-ai-proves-its-worth-but-few-scale-impact# (accessed Oct. 15, 2020)

[9] K. Kalafala et al., "A Collaborative Data Model for AI/ML in EDA", Si2 Digital Publications, Aug. 2020. https://si2.org/product/collaborative-data-model (accessed Oct. 19, 2020).

[10] Coursera. https://www.coursera.org/ (accessed Oct. 13, 2020).

